

Consensus backbones for directed economic and financial networks: sovereign risk spillovers and the emerging-market core

Jorge Antonio Chan-Lau

jchanlau@gmail.com

Frontier Quant Analytics

Research Article

Keywords: directed networks, arborescence, consensus backbone, network sparsification, variance decomposition, sovereign risk spillovers, credit default swaps, emerging markets

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9708604/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Consensus Backbones for Directed Economic and Financial Networks: Sovereign Risk Spillovers and the Emerging-Market Core

May 2026

Author Details

Jorge A. Chan-Lau

Frontier Quant Analytics LLC

NUS AIDF Credit Research Initiative

Corresponding author: Jorge A. Chan-Lau

Email: jchanlau@gmail.com

Corresponding author postal address: McLean, VA 22101, USA

Acknowledgements

None

Declaration of Competing Interest

The author declares that there are no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Author Contributions

Jorge A. Chan-Lau: Conceptualization, methodology, software, formal analysis, investigation, data curation, writing–original draft, writing–review and editing, visualization.

Abstract

Consensus backbones provide a directed dimensionality-reduction layer for dense weighted networks by aggregating multiple maximum arborescences into a sparse and interpretable union. This paper develops the method for directed economic and financial networks and uses a 36-country Diebold–Yilmaz spillover network as an illustration. The resulting backbone retains 3.8% of total network weight using only 2.1% of edges, revealing the dominant hierarchical spillover structure in the estimated dependence graph. A striking empirical finding emerges from the application: under scale-invariant Sinkhorn normalization, all 21 advanced economies in the sample are absent from the backbone, while 15 emerging market and commodity-linked sovereigns constitute the entire active spillover hierarchy. This pattern reflects a fundamental asymmetry in sovereign risk transmission: advanced-economy spillovers are diffuse and nearly symmetric, whereas emerging market spillovers are concentrated along specific bilateral channels that the backbone isolates as the dominant directional structure. The method identifies key structural transmitters, receivers, and bridge economies in the illustrative example, while preserving the possibility of alternative spillover routes that a single tree would miss.

Keywords: directed networks, arborescence, consensus backbone, network sparsification, variance decomposition, sovereign risk spillovers, credit default swaps, emerging markets.

JEL codes: C32, C63, D85, F34, G01, G15.

1 Introduction

Economic and financial systems exhibit intricate webs of interdependence—cross-border exposures, production linkages, trade flows, and payment obligations—that allow disturbances to propagate through multiple, often overlapping, channels. Mapping and simplifying these dense networks has become essential for understanding contagion, systemic risk, and the topology of economic influence.

Early approaches used the Minimum Spanning Tree (MST) to filter correlation networks into sparse, interpretable structures that preserve the most significant linkages (Mantegna 1999). The MST extracts a backbone, or the minimal subgraph capturing the essential structure of dependencies (Esmalifalak 2021; Valle and Urbina 2022). However, MSTs and related filtering techniques operate on undirected networks, ignoring the asymmetry that characterizes many economic relationships, such as creditor–debtor, exporter–importer, or parent–subsidiary hierarchies.

Directed analogues known as arborescences or branchings offer a natural extension: a rooted, directed spanning tree in which all edges emanate from a source node, tracing directional dependence or spillovers through the system characterized by the network (Coscia 2018; Edmonds 1967). Arborescences thus represent directed backbones, emphasizing hierarchical propagation rather than symmetric co-movement.

Yet, an arborescence is a single-tree representation which may be overly restrictive. Financial and economic networks display redundant paths, that is, alternative channels through which shocks or spillovers can travel. The branching-pack (or k -arborescence) concept from combinatorial optimization generalizes this by identifying multiple, edge-disjoint arborescences rooted at one or more sources (Bérczi and Frank 2016). The union of these directed trees forms a consensus backbone: a sparse yet resilient structure that preserves alternative high-weight spillover routes without overwhelming density.

This paper proposes the consensus backbone as a natural way to analyze weighted directed economic and financial networks. Our framework constructs this structure as the union of multiple high-weight arborescences, or branching packs, each representing a maximal directional hierarchy rooted at a key node. This framework generalizes traditional backbone filtering methods, such as minimum spanning trees (MSTs) and statistical filters, by incorporating both directionality and redundancy. It allows us to identify the dominant spillover paths conditional on the estimated model, while achieving an interpretable and structurally coherent representation of the predictive dependence structure of the network.

To our knowledge, this is the first application of directed consensus backbones to economic and financial networks. This contribution bridges three previously separate research streams: backbone extraction in finance, which has relied on undirected methods such as MSTs, planar maximally filtered graphs, and statistical filters (Mantegna 1999; Esmalifalak 2021; Valle and Urbina 2022; Yassin et al. 2023); arborescence-based analysis in network science, which has been confined to hierarchical characterization of directed graphs (Coscia 2018) and combinatorial optimization (Bérczi and Frank 2016); and the Diebold–Yilmaz variance-decomposition framework

(Diebold and Yilmaz 2012), which generates dense directed networks but whose analysis has been limited to scalar connectedness indices. By bringing these streams together, the consensus backbone provides a principled method for extracting sparse, directed, and interpretable structures from dense spillover networks.

To illustrate its application, we construct a global spillover network based on the standard generalized forecast error variance decomposition of a multivariate vector autoregression (VAR) following Diebold and Yilmaz (2012) and Pesaran and Shin (1998). While the variance decomposition network is dense and high-dimensional, with each pair of economies connected by bidirectional edges representing risk contributions or spillovers, the consensus backbone reduces it to a sparse, hierarchical structure that reveals the dominant predictive dependence structure of sovereign risk spillovers conditional on the estimated VAR.

The remainder of the paper is organized as follows. Section 2 presents a brief overview of backbones, how they relate to earlier analysis of economic and financial networks, and briefly discusses why consensus backbones are better suited to analyze directed networks. Section 3 offers a conceptual overview of consensus backbones followed by a concise discussion of the different structures involved in their construction, including the Chu–Liu/Edmonds algorithm; branching packs; the use of matrix normalizations to account for scale effects; and a small worked example illustrating the construction step by step. The methods described in this section are put to use in Section 4, where spillovers in a variance decomposition network are summarized using consensus backbones. Section 5 discusses the implications of the key empirical finding—the absence of advanced economies from the backbone—for the understanding of sovereign risk transmission. Section 6 concludes and discusses extensions.

2 Literature review

Network analysis has become a central tool for studying interconnectedness in financial and macroeconomic systems. Surveys such as Emmert-Streib et al. (2018) and Bougheas and Kirman (2014) review the evolution of financial-network research from early correlation matrices to multilayer representations of macro-financial linkages. These studies emphasize that the network topology, characterized by properties such as density, clustering, and hierarchy, shapes systemic outcomes such as contagion, liquidity shortages, and efficiency of risk sharing.

The concept of a backbone, a reduced subgraph retaining only the most significant edges, appears throughout the finance literature. A substantial number of studies rely on minimal spanning trees (MSTs) and related filtering methods to extract backbone structures from dense correlation matrices. These approaches yield undirected, cycle-free representations that preserve the strongest bilateral relations and ensure global connectivity. Such backbones have been used to analyze the hierarchical structure of stock markets (Mantegna 1999), cross-border exposures (Spelta and Araújo 2011), corporate control (Vitali, Glattfelder, and Battiston 2011), payment and settlement networks (Bech and Sørensen 2018; Gaffeo and Molinari 2016), trade and production networks (Mastrandrea et al. 2014), and the flow of goods and credit (Huremovic et al. 2020). More broadly, Esmalifalak (2021) summarizes several correlation-network methodologies (MST, planar maximally filtered graphs, disparity filters), emphasizing their role as undirected backbones of

financial systems. Whether structural (preserving topological features) or statistical (filtering edges relative to a null model, as in Valle and Urbina (2022)), these methods share a common limitation: they treat linkages as symmetric, discarding directional information.

Yet many forms of economic linkage, such as credit exposures, trade, and risk spillovers, are inherently directional. A directed network is better represented by an arborescence, a directed spanning tree that captures spillover paths from a root to all other nodes. An arborescence thus serves as a directed backbone, retaining the strongest and most interpretable directional paths, analogous to how MSTs retain the strongest symmetric dependencies. By replacing the undirected minimal tree with a directed maximal branching, one transforms the backbone concept from “minimal connectivity” to “maximal directional influence.”

A limitation of both MSTs and single arborescences is that they may yield overly brittle representations: a single path failure can sever parts of the network. The branching-pack (or k -arborescence) framework of Bérczi and Frank (2016) addresses this by constructing k edge-disjoint arborescences, all rooted at the same node, that together maximize total edge weight. This approach extends the Chu–Liu/Edmonds algorithm (Chu and Liu 1965; Edmonds 1967) and generates a multi-path directed backbone from a given root’s perspective, which can capture redundant spillover channels analogous to alternative liquidity routes in interbank systems. The union of these k arborescences forms a richer yet still sparse subgraph, one that preserves multiple high-weight spillover routes while remaining resilient to single-path failures.

While k -arborescences provide robustness within a single rooted perspective, a further level of robustness can be achieved by aggregating across multiple perspectives or methods. This leads to the concept of a consensus backbone: a structure obtained by combining multiple backbone representations, whether derived from different root nodes, different extraction algorithms, or different filtering criteria, and retaining edges that appear consistently across them. Edges that survive this aggregation are deemed structurally robust. Yassin et al. (2023) propose a methodological framework for comparing backbone extraction techniques and explicitly computing their consensus intersection. Similarly, Grady, Thiemann, and Brockmann (2012) define link salience, constructing a “salience skeleton” representing the consensus over all nodes’ shortest-path trees. Related concepts include the Union of Shortest-Path Trees, applied to brain networks (Meier, Tewarie, and Van Mieghem 2015), and the ultrametric backbone of Rozum and Rocha (2024), obtained as the union of minimum-spanning forests.

This arborescence-based perspective also connects the proposed framework to related work on hierarchical and directional economic networks. Consensus backbones can be viewed as generalizing methods that use simpler directed structures, including the arborescence-based hierarchy index for directed networks (Coscia 2018), work on the theoretical foundations of economic directionality (Elliott, Golub, and Jackson 2019), and contagion models emphasizing stress propagation along edges capturing direct exposures (Serri et al. 2016). In a macro-financial risk context, the findings of Joseph et al. (2014) and Sharma, Chakrabarti, and Chakraborti (2018) suggest that global interdependencies often form layered or tree-like structures, making arborescent summaries a natural representation. The same logic is relevant for policy-oriented analysis of economic networks, such as the ESRB’s work on network reconstruction (Anand

et al. 2018) and the DebtRank approach to systemic importance (Battiston et al. 2012), where sparse and interpretable structures can help preserve the critical interconnections of dense exposure or dependence networks.

Despite these related strands, applications of consensus backbones to economics and finance networks appear to be absent. Existing studies extract a single backbone (typically an MST, PMFG, or statistical filter) but do not combine multiple backbones or directed trees into a unified structure. This paper addresses that gap by constructing consensus backbones from k -arborescences rooted at different key nodes. The branching-pack framework provides robustness within each rooted perspective by generating multiple edge-disjoint spillover trees; aggregating across roots then provides robustness across perspectives, yielding a directed backbone that is both economically interpretable and resilient to the choice of reference node.

3 Consensus backbones

3.1 Conceptual overview

The method relies on the graph representation of a network and the construction of maximum directed arborescences, branching packs, and the corresponding consensus backbone. A directed arborescence is a spanning tree with a root node, r , with in-degree 0 at r , and in-degree 1 for all other nodes, all reachable from the root node (Chu and Liu 1965; Edmonds 1967). The maximum arborescence maximizes the sum of the selected edge weights subject to acyclicity and in-degree constraints. It is the directed analogue of the minimum/maximum spanning tree of an undirected graph (Tarjan 1977; Wang et al. 2023).

Each maximum arborescence gives a dominant path structure from the root node. To capture alternative paths, multiple diversified arborescences are computed across multiple roots and multiple runs per root. The collection of these disjoint arborescences constitutes the branching pack (Bérczi and Frank 2016). Finally, the consensus backbone is the unified graph structure obtained by merging the branching pack arborescences (Bérczi and Frank 2016; Georgiadis 2003). It preserves the most significant directed edges that recur across them, providing a sparse yet comprehensive representation of the dominant spillover paths of the original graph. Each construct is discussed in detail below.

3.2 Normalization and relative arborescence

In a directed graph, the edge weight, $w_{i,j}$, measures the influence of node i on j . One way to measure the total influence of node i on the graph, or its importance as an emitter, is to calculate its out-strength, $\sum_{j \neq i} w_{i,j}$. Similarly, the total influence of the graph on node i , or its importance as a receiver, can be measured by its in-strength, $\sum_{j \neq i} w_{j,i}$. The total strength of the node is the sum of its in-strength and out-strength.

A node’s influence, whether as an emitter or a receiver, is determined by its scale: a node connected by a few high-weight edges can exert as much influence as one connected by many low-weight edges. Without proper scaling, these disparate connection profiles can create significant distortions in the analysis of the graph. In a raw graph, absent scaling, high-magnitude weights

often overshadow the structural importance of topological variety, leading to a bias where a single “heavy” connection carries more weight than a broad “community” of interactions. Scaling mitigates this by normalizing these values, ensuring that influence is measured relative to the global distribution of the graph rather than just the absolute values of individual edges.

Normalization of the adjacency matrix $W = [w_{ij}]$ removes scaling biases, equalizing the impact of nodes regardless of whether they possess a few high-weight edges or many low-weight ones. The choice of normalization represents a critical methodological decision, as each approach highlights a distinct facet of network dynamics. Three common techniques are employed based on the specific analytical objective:

Row-stochastic normalization ($\tilde{w}_{ij} = w_{ij} / \sum_k w_{ik}$) ensures that each node’s outgoing weights sum to unity, emphasizing the relative distribution of influence emitted by each source node independent of its total out-strength. This approach evaluates *emission profiles* and is particularly effective for uncovering dominant transmitters of risk who distribute influence with high asymmetry.

Column-stochastic normalization ($\tilde{w}_{ij} = w_{ij} / \sum_k w_{kj}$) standardizes total in-strength to unity, focusing on the relative exposure of each receiver node to its various sources. This is used to assess *receiver vulnerability*, highlighting target nodes that are disproportionately dependent on specific sources for exposure and susceptibility analysis.

Doubly stochastic normalization ($\tilde{W} = D_r^{-1} W D_c^{-1}$) uses the Sinkhorn–Knopp adjustment to iteratively balance both inflows and outflows. By controlling for node-specific scale effects, it isolates the *latent spillover structure* and reveals edges that represent exceptionally strong linkages. This scale-invariant perspective is the most neutral and is recommended as the default for exploratory analysis (Sinkhorn 1967; Knight 2008; Idel 2016; Chakrabarty and Khanna 2018).

Once the adjacency matrix is properly scaled, application of the Chu–Liu/Edmonds algorithm identifies *relative arborescences*, which, when assembled into a consensus backbone, recover the directional pathways in the graph. This approach is particularly useful in the analysis of economic and financial systems characterized by scale asymmetries, such as cross-border risk spillovers, trade dependencies, and financial contagion. The algorithm and the consensus backbones are explained next.

3.3 Chu–Liu/Edmonds algorithm

The Chu–Liu/Edmonds algorithm (Chu and Liu 1965; Edmonds 1967) identifies a *maximum* (or *minimum*) *spanning arborescence* in a directed, weighted graph rooted at a designated node r . An arborescence is a directed tree spanning all reachable nodes such that exactly one directed path exists from r to every other node—the directed analogue of an undirected spanning tree. Whether one seeks a minimum or maximum arborescence depends on the semantics of edge weights. When weights represent costs or frictions, as in credit intermediation chains, a minimum arborescence identifies the lowest-cost funding paths from a central lender to all borrowers. In

contrast, when weights represent exposure or influence, as in risk spillover networks similar to the illustrative example below, a maximum arborescence identifies the stronger spillover paths.

Without loss of generality, the discussion hereafter concerns the maximum arborescence. Formally, given a directed graph $G = (V, E)$ with edge weights w_{ij} for edges $(i, j) \in E$, the maximum arborescence problem can be written as:

$$\begin{aligned}
& \max_{x_{ij}} \quad \sum_{(i,j) \in E} w_{ij} x_{ij} \\
& \text{s.t.} \quad \sum_{i:(i,j) \in E} x_{ij} = 1, \quad \forall j \in V \setminus \{r\}, \\
& \quad \sum_{i:(i,j) \in E} x_{ij} = 0, \quad \text{for } j = r, \\
& \quad x_{ij} \in \{0, 1\}, \\
& \quad \text{and no directed cycles are formed.}
\end{aligned} \tag{1}$$

The solution $x_{ij} = 1$ identifies the set of directed edges forming the arborescence T_r rooted at r . Algorithm 1 describes the Chu–Liu/Edmonds procedure.

Algorithm 1: Chu–Liu/Edmonds Maximum Arborescence

Input: Directed graph $G = (V, E)$, edge weights w_{ij} , root node r

Output: Maximum spanning arborescence T_r rooted at r

foreach node $j \in V \setminus \{r\}$ **do**

 Select incoming edge $(i^*, j) \leftarrow \arg \max_{i:(i,j) \in E} w_{ij}$;

if selected edges form a directed cycle C **then**

 Contract C into a super-node v_C ;
 Adjust edge weights: $w'_{iv_C} \leftarrow w_{ij} - w_{i^*j} + w_C^*$ for edges entering C ;
 Recurse on the contracted graph G' ;

Expand contracted nodes to recover T_r ;

return T_r

3.4 Branching packs and consensus backbones

Although a single arborescence T_r captures the most dominant propagation structure from the root, it may be too sparse to represent systems where multiple plausible routes coexist. A *branching pack*—as defined in Bérczi and Frank (2016)—is a collection of k edge-disjoint arborescences $\{T_r^{(1)}, T_r^{(2)}, \dots, T_r^{(k)}\}$ rooted at the same node r . Each arborescence is obtained by solving a Chu–Liu/Edmonds maximum branching problem, with the constraint that its edges do not overlap with those selected in previous solutions. This iterative optimization produces a family of arborescences that are all optimal with respect to their residual graphs, thereby capturing distinct high-weight propagation routes emanating from the same source.

The union of these k arborescences forms a *consensus backbone*, preserving multiple high-weight spillover paths while maintaining interpretability and bounded density. In economic networks, such a structure enables analysts to visualize alternative channels of spillovers without resorting to arbitrary edge-weight thresholds, which could prove useful when analyzing competing routes

of contagion, influence, or trade. The consensus backbone thus balances *parsimony* (through directed-tree constraints) with *representativeness* (via diversified inclusion of key connections).

A natural question is how to select k . Edmonds' branching theorem (Edmonds 1967) establishes that the maximum number of edge-disjoint arborescences rooted at r equals the minimum in-degree among all reachable nodes:

$$k_{\max} = \min_{v \neq r} \deg^-(v), \quad (2)$$

where $\deg^-(v)$ is the in-degree of node v . This provides a structural upper bound on the size of any branching pack. In practice, however, k is determined by modeling and interpretive trade-offs rather than set at its theoretical maximum. The goal is to select k such that the consensus backbone:

- retains a substantive share of total network edge weight (weight ratio),
- connects a sufficient fraction of ordered node pairs by a directed path (reachability),
- avoids excessive edge density (density), and
- faithfully reflects the node-importance rankings of the full network, as measured by the R^2 between backbone out-strength and full-network out-strength (strength preservation).

These four properties are jointly tracked as k increases, and k^* is selected at the point of diminishing returns: the value beyond which the marginal improvement in a composite of these diagnostics falls below a pre-specified threshold (Tarjan 1977; Camerini, Fratta, and Maffioli 1979; Bérczi and Frank 2016).

Specifically, we adopt a *weighted composite criterion* that balances all four diagnostics simultaneously. Each diagnostic is normalized to $[0, 1]$ by dividing by its maximum across k , and a composite score is formed as the equal-weight average:

$$S(k) = \frac{1}{4} [\widetilde{W}(k) + \widetilde{D}(k) + \widetilde{R}(k) + \widetilde{Q}(k)], \quad (3)$$

where $\widetilde{W}$, $\widetilde{D}$, $\widetilde{R}$, and $\widetilde{Q}$ denote the normalized weight ratio, density, reachability, and out-strength R^2 , respectively. The selected k^* is the largest k whose marginal composite gain $\Delta S(k) = S(k) - S(k-1)$ exceeds 10% of the first-step gain $\Delta S(2)$:

$$k^* = \max \{ k \geq 2 : \frac{\Delta S(k)}{\Delta S(2)} \geq 0.10 \}. \quad (4)$$

Having defined the logic for multipath aggregation and the criteria for selecting the optimal number of arborescences (k^*), the complete construction process is consolidated in Algorithm 2.

3.5 Consensus backbone algorithm

Algorithm 2 summarizes the consensus backbone construction. The inputs are the weighted adjacency matrix W , the number of roots k_r , and the number of arborescences per root k_p . The choice of k_p involves a parsimony tradeoff: too few arborescences yield an overly sparse

backbone that may miss important alternative paths; too many reduce parsimony without adding structural information as successive arborescences capture increasingly weak or redundant links. In the application, k_p is selected data-adaptively using a composite diagnostic criterion; the data-selected value is denoted k^* (see Section 4.3). The key design choices are:

- **Root selection.** Roots are chosen as the top- k_r nodes by out-strength in the (possibly normalized) adjacency matrix $\tilde{W}$. This endogenous selection ensures that the dominant spillover sources serve as starting points for the arborescences.
- **Edge-disjointness.** Within each root’s pack of arborescences, previously selected edges are penalized (not removed) by a compound multiplicative factor $\gamma^{c_{uv}}$, where $\gamma \in (0, 1)$ is the penalty parameter and c_{uv} counts how many times edge (u, v) has been used. This encourages diversity without making subsequent arborescences infeasible when the graph has limited edge-disjoint structure. The penalty is applied *within* each root, not across roots.
- **Reachable arborescences.** The Chu–Liu/Edmonds algorithm returns the maximum arborescence over the subgraph reachable from the specified root. If some nodes are unreachable from a given root, they will not appear in that arborescence. Consequently, the final union need not span every node—some countries may remain isolated if they are not reached by any root’s arborescence.
- **Frequency scoring.** Each edge in the union receives a frequency score f_{uv} equal to the number of arborescences (across all roots) in which it appears. Higher-frequency edges represent the most robust spillover links.

Algorithm 2: Consensus Backbone Construction

Input: Adjacency matrix $W \in \mathbb{R}_{\geq 0}^{N \times N}$, normalization mode, k_r (roots), k_p (arborescences per root), penalty $\gamma \in (0, 1)$

Output: Backbone graph $B = (V_B, E_B)$ with edge weights and frequency scores


```
 $\tilde{W} \leftarrow \text{NORMALIZE}(W, \text{mode});$   
 $G \leftarrow \text{BUILDDIGRAPH}(\tilde{W}) ;$  // remove diagonal; edge weight =  $\tilde{w}_{ij}$   
 $s_i \leftarrow \sum_j \tilde{w}_{ij}$  for all  $i$  ; // out-strength  
 $\mathcal{R} \leftarrow \text{top-}k_r \text{ nodes by } s_i ;$  // root set  
 $\mathcal{T} \leftarrow \emptyset ;$  // collection of arborescences  
foreach  $r \in \mathcal{R}$  do  
|  $c_{uv} \leftarrow 0$  for all  $(u, v) \in E(G) ;$  // reset usage counts per root  
| for  $j = 1, \dots, k_p$  do  
| |  $G_j \leftarrow G$  with  $w_{uv} \leftarrow w_{uv} \cdot \gamma^{c_{uv}}$  for all  $(u, v)$ ;  
| |  $T \leftarrow \text{CHULIU/EDMONDS}(G_j, r) ;$  // max arborescence over reachable  
| | subgraph  
| |  $\mathcal{T} \leftarrow \mathcal{T} \cup \{T\};$   
| | Update  $c_{uv} \leftarrow c_{uv} + 1$  for each  $(u, v) \in E(T)$ ;  
| end  
end  
 $B \leftarrow \bigcup_{T \in \mathcal{T}} T ;$  // union of all arborescences  
foreach  $(u, v) \in E(B)$  do  
|  $w_{uv}^B \leftarrow \max_{T \ni (u, v)} \tilde{w}_{uv}^{(T)} ;$   
|  $f_{uv} \leftarrow |\{T \in \mathcal{T} : (u, v) \in E(T)\}| ;$  // frequency score  
end  
return  $B$  with weights  $\{w_{uv}^B\}$  and scores  $\{f_{uv}\};$ 
```

3.6 A worked example

To make the construction concrete, consider a six-node directed network. This stylized economy has directed risk spillovers. Figure 1 displays the full network with edge weights and Figures 2 to 4 illustrates the several steps involved in constructing the consensus backbone.

Figure 1: Full 6-node directed network (toy example).

A complete directed network on six nodes $\{A, B, C, D, E, F\}$ with illustrative edge weights representing pairwise risk spillovers. Edge widths are proportional to weight.



Source: Author's illustration.


Step 1: Single maximum arborescence. Applying the Chu–Liu/Edmonds algorithm with root A yields the maximum arborescence $T_A^{(1)}$, shown in Figure 2. This tree spans all nodes via the highest-weight directed paths from A , capturing the single most dominant spillover hierarchy. The tree, however, may miss important alternative routes.

Step 2: Second arborescence (branching pack with $k = 2$). After removing the edges used in $T_A^{(1)}$, a second maximum arborescence $T_A^{(2)}$ is computed from the residual graph, as shown in Figure 3. This tree captures the next-best alternative spillover routes.

Step 3: Consensus backbone. The consensus backbone is the union $B = T_A^{(1)} \cup T_A^{(2)}$, shown in Figure 4. Edges that appear in both arborescences are assigned higher frequency scores and represent the most robust spillover channels. The backbone is sparser than the full network but richer than any single arborescence, preserving the key alternative paths.

Figure 2: First maximum arborescence $T_A^{(1)}$ (toy example).


Maximum arborescence $T_A^{(1)}$ rooted at node A is the single directed spanning tree that maximizes total edge weight and captures the dominant spillover paths.



Source: Author's illustration.

Figure 3: Second maximum arborescence $T_A^{(2)}$ (toy example).


Second maximum arborescence $T_A^{(2)}$, computed from the residual graph after removing edges in $T_A^{(1)}$. Together, $\{T_A^{(1)}, T_A^{(2)}\}$ form a branching pack of size $k = 2$.



Source: Author's illustration.

Figure 4: Consensus backbone from $k = 2$ branching pack (toy example).

The consensus backbone formed as the union of the two arborescences. Edge thickness reflects frequency (number of arborescences selecting the edge). The backbone retains the dominant and alternative spillover paths while discarding the weakest connections.



Source: Author's illustration.

This example illustrates the key intuition: a single arborescence reveals the dominant hierarchy, the branching pack uncovers alternative routes, and the consensus backbone synthesizes both into a sparse, interpretable directed structure. In the application that follows, the same logic is applied to a 36-country network with over a thousand edges, where the gains from sparsification are substantial.

4 Application: sovereign directed spillover network

We illustrate the use of consensus backbones to identify dominant spillover paths in a sovereign spillover graph $G = (V, E)$, where V denotes the set of countries (nodes), E is the set of edges (i, j) with weight w_{ij} equal to the risk spillover from country i to country j , which can be represented by a matrix W .

4.1 Sovereign risk measure and credit default swap data

The risk measures in the analysis are the weekly changes of the logit transformation of the 5-year probability of default (PD) implied from USD-denominated 5-year credit default swap spreads, S . The spreads are converted to probabilities of default assuming a recovery rate, R , of 60 percent and a flat term structure of default risk according to:

$$PD_T = 1 - \exp\left(-\frac{S \times T}{1 - R}\right), \quad (5)$$

which yields the risk measure, $\text{logit}(PD_T)$:

$$\text{logit}(PD_T) = \ln \left(\frac{PD_T}{1 - PD_T} \right), \quad (6)$$

where T is the 5-year time horizon. The standard logit maps a probability to the log-odds of default. Because the analysis is based on variance decompositions of the first differences of $\text{logit}(PD_T)$, the sign of the level transformation does not affect the GFEVD or any downstream backbone calculations.

The CDS spreads used to compute the logit risk measures are sourced from LSEG Datastream, which reports mid-prices collected by CMA Datavision. The data are sampled at daily frequency and cover the period December 15, 2008 to May 27, 2025. For the analysis, the data are resampled at weekly frequency, with each weekly data point representing the end-of-week value, for a total of 870 observations. The sample covers 36 economies: Australia, Austria, Belgium, Brazil, Chile, China, Colombia, Denmark, France, Germany, Hong Kong, Hungary, Indonesia, Ireland, Israel, Italy, Japan, Korea, Malaysia, Mexico, the Netherlands, New Zealand, Norway, Panama, Peru, the Philippines, Poland, Portugal, South Africa, Spain, Sweden, Thailand, Turkey, the United Kingdom, the United States, and Vietnam.

4.2 Spillover network construction

The illustrative network, as of end-May 2025, is the variance-decomposition network first proposed by Diebold and Yilmaz (2012) to measure directional volatility spillovers. In the Diebold–Yilmaz interpretation, a spillover is the fraction of one variable’s forecast error variance attributable to shocks from other variables in the system. To construct this network we first estimate a vector autoregression (VAR) for the first difference of the logit PDs to capture their dynamic interrelationship. We adopt a VAR(8) specification estimated by ordinary least squares (OLS), corresponding to eight weeks of lag structure. This choice follows the convention in the weekly Diebold–Yilmaz literature and balances two competing concerns: enough lags to capture plausible short-to-medium-run dynamics (two months), while maintaining adequate degrees of freedom ($T - p = 535$ observations, $535 - (36 \times 8 + 1) = 246$ degrees of freedom per equation). Standard information criteria applied to this system select lag 1, an artifact of the severe dimensionality penalty in a 36-variable system with roughly 13 feasible lags; the AIC selects the feasibility ceiling of 13 lags, confirming that the data contain predictability at longer horizons. Section 4.7.1 reports a systematic sensitivity analysis across lag lengths of 1, 4, 8, and 12, showing that the VAR(8) backbone is broadly representative of the intermediate-lag estimates.

We then compute the generalized forecast error variance decomposition (GFEVD) of Pesaran and Shin (1998), which is invariant to the order of the variables, to measure cross-country contributions to forecast error variance. The GFEVD measures predictive variance shares, not structural causality. An entry θ_{ij} is the share of country i ’s forecast error variance at horizon H explained by innovations in country j ’s reduced-form VAR residuals. Such predictive shares may reflect common factors, global shocks, or specification error, not just bilateral spillovers. The consensus backbone should therefore be read as a *predictive dependence skeleton*: a sparse summary of the

strongest directional forecast links, not a map of structural causation. Accordingly, “spillover path” and “spillover hierarchy” are used throughout in this predictive, not causal, sense.

The GFEVD yields a dense pairwise variance contribution adjacency matrix, with all off-diagonal entries strictly positive. After excluding each variable’s own contribution to its forecast error variance (the diagonal entries), the illustrative network is a fully connected directed graph with $N(N - 1) = 36 \times 35 = 1,260$ edges. The high density and dimensionality of the spillover network is precisely the challenge that motivates the consensus backbone approach: without sparsification, it is difficult to distinguish the dominant spillover paths from the noise induced by weak bilateral linkages.

Finally, the computation of the GFEVD network requires specifying the time horizon as spillovers are not instantaneous and evolve over time. For the purpose of this analysis, which serves mainly to illustrate the methodology, the horizon is set equal to 52 weeks, or one year, which we consider more appropriate from a macro-prudential perspective. To assess the robustness of the results, a sensitivity analysis is performed for horizon values of 4, 12, and 26 weeks (see subsection 4.7.3).

4.3 Selection of k-star


Applying the composite selection criterion defined in Section 3.4, we determine the optimal number of arborescences k^* by evaluating the trade-offs between sparsification and information retention. Figure 5 illustrates the evolution of the four underlying diagnostic metrics—union weight ratio, density, reachability, and out-strength R^2 —as k increases.

In the baseline OLS VAR(8) network, the composite score $S(k)$ identifies $k^* = 3$ as the optimal stopping point. At this value, the backbone captures 99.5% of the maximum composite score. More importantly, the marginal gain of moving from $k = 3$ to $k = 4$ represents only 1.1% of the initial gain from the first arborescence, falling well below the 10% threshold for inclusion.

As shown in the diagnostic curves, the choice of $k^* = 3$ is driven by rapid structural saturation. Reachability plateaus at $k = 2$, indicating that the union of just two arborescences is sufficient to cover the reachable subgraph from each root. Furthermore, the correlation between backbone and original out-strengths (R^2) reaches its maximum at $k = 3$. Beyond this point, adding further arborescences increases density without providing meaningful gains in strength preservation or topological coverage. Consequently, three arborescences provide a parsimonious representation of the system’s dominant spillover hierarchy.

Figure 5: Diagnostic curves for k^* selection (OLS VAR(8) baseline)

The figure plots the four diagnostic metrics used to calculate the composite selection score $S(k)$ for the Sinkhorn-normalized network. Top left: cumulative weight ratio. Top right: backbone density. Bottom left: reachability (fraction of reachable node pairs). Bottom right: R^2 between backbone and original out-strength. The vertical dashed line at $k^* = 3$ indicates the optimal number of arborescences selected by the composite criterion (4).



Source: Author's calculations.


4.4 Sparsification

A variance decomposition network is a large, fully connected graph with a total of $N(N - 1)$ edges, where N is the number of nodes, once variables' own contributions to their error variances are excluded. In the illustrative network, the number of directional edges among countries is 1,260. Because of the high density and dimensionality of the network, the analysis of variance decomposition networks has been mainly confined to examining several connectedness measures based on the average and total error variance contributions of the nodes (Gabauer, Chatziantoniou, and Stenfors 2023).

The consensus backbone dramatically reduces this complexity. Figure 6 presents side-by-side heatmaps of the full Diebold–Yilmaz adjacency matrix and the consensus backbone adjacency matrix, visually illustrating the sparsification achieved.

Figure 6: Full variance decomposition network vs. consensus backbone (heatmaps).

Left panel: full 36×36 variance decomposition adjacency matrix (excluding diagonal). Right panel: consensus backbone adjacency matrix. Countries are ordered identically in both panels. The backbone retains the dominant structure while eliminating weak linkages.



Source: Author's calculations.

Table 1 summarizes the sparsification statistics under four normalization modes. Although the backbone retains only a modest share of total network weight, consistent with its role as a structural sparsifier rather than a weight maximizer, it achieves this with a small fraction of edges while preserving directed connectivity, a property that naive thresholding at comparable density does not guarantee.

Table 1: Sparsification statistics by normalization mode.

Normalization	Edges (original)	Edges (backbone)	Weight coverage (%)	Density
Raw	1,260	177	27.3	0.1405
Row-stochastic	1,260	56	9.7	0.0444
Column-stochastic	1,260	76	12.4	0.0603
Sinkhorn	1,260	27	3.8	0.0214

Notes: Edges (original) is $N(N - 1) = 1,260$ in all cases. Edges (backbone) is the number of edges in the consensus backbone at the selected $k^* = 3$. Weight coverage is the ratio of total backbone edge weight to total original edge weight. Density is the ratio of backbone edges to the maximum possible number of edges. All backbones use the OLS VAR(8) baseline GFEVD, Sinkhorn normalization is the primary specification used throughout the application. The raw backbone is notably denser because unscaled GFEVD entries vary substantially across country pairs, allowing the arborescence algorithm to stack up edges along high-weight rows; normalization removes this scale heterogeneity and concentrates the backbone on structurally dominant paths.


Source: Author's calculations.

Beyond the edge-count statistics, Figure 7 shows how the backbone *topology* responds to the normalization choice. Core spillover paths are preserved across all four modes; peripheral edges

vary with the scale adjustment, reflecting the different emphasis each normalization places on source-side versus target-side magnitudes.

Figure 7: Sensitivity of the consensus backbone to normalization.

Consensus backbones under four normalization modes (OLS VAR(8) baseline). Top left: raw weights (177 edges). Top right: row-stochastic (56 edges). Bottom left: column-stochastic (76 edges). Bottom right: Sinkhorn (27 edges). The backbone is substantially sparser under normalizations that remove cross-country scale differences, with Sinkhorn yielding the most parsimonious structure by jointly controlling source- and target-side magnitudes.



Source: Author's calculations.

4.5 Spillover hierarchy

Two substantive points follow from these sparsification results. First, the low weight coverage of the Sinkhorn backbone should not be read as poor information retention. The objective of the consensus backbone is not to maximize the sum of retained edge weights, but to preserve a coherent set of dominant directed paths under arborescence constraints. A threshold rule can retain more aggregate weight by selecting many locally large but mutually disconnected edges; the backbone instead sacrifices some raw weight coverage in order to retain a navigable spillover

structure. Second, the contrast between the raw and normalized backbones shows that scale adjustment materially affects sparsity. Without normalization, large absolute GFEVD entries allow the algorithm to stack edges along high-weight rows, producing a much denser graph. The row, column, and Sinkhorn normalizations remove different forms of scale heterogeneity and therefore concentrate the backbone on fewer structurally dominant paths. Section 4.5 develops the substantive interpretation of the Sinkhorn-normalized backbone.


We employ the Sinkhorn-normalized backbone as our primary specification because it simultaneously accounts for source- and target-side scale effects. By prioritizing relative spillover intensity over absolute magnitude, this normalization reveals the underlying structural dependencies between country pairs. Consequently, prominence within this backbone implies high relative importance *after controlling for scale*, which is not necessarily synonymous with being a top absolute contributor to global sovereign risk. For instance, a country with moderate aggregate spillovers, but with high concentration along specific bilateral channels, may appear more central in the Sinkhorn backbone than its raw DY connectedness index would suggest.

The scale-invariant weights provide the foundation for uncovering the network’s latent architecture. The resulting consensus backbone reveals a clear hierarchical dependence structure, which we visualize in Figure 8 using a Kamada–Kawai hierarchical layout. This rendering intuitively positions root nodes, those with the highest relative out-strength, at the top of the system, while cascading down to peripheral nodes at the base. Complementing the hierarchical view, Figure 9 employs a radial layout centered on the node with the highest relative out-strength. This rendering intuitively illustrates how spillovers radiate outward from the primary transmitter to the rest of the network and emphasizes the reach of the dominant root and the transition from the core to the most isolated peripheral economies.

Finally, table 2 reports node-level statistics for all 36 countries, including degree and strength measures as well as a backbone centrality score measuring each node’s role as a non-leaf transmitter across arborescences. Because the consensus backbone is the union of arborescences computed over the subgraph *reachable* from each selected root rather than spanning arborescences that cover every node by construction some countries may appear with zero in-degree and zero out-degree if they are not reached by any root’s arborescence.

Figure 8: Consensus backbone: hierarchical spillover structure.


Consensus backbone of the sovereign spillover network (OLS VAR(8), Sinkhorn normalization, $k = k^* = 3$). Edge color intensity reflects backbone weight. The layout reveals the dominant predictive spillover paths from the nodes with highest out-strength to peripheral receivers. Node prominence reflects scale-adjusted relative centrality in the Sinkhorn-normalized network; see Table 3 for the formal quartile-based role classification.



Source: Author's calculations.

Figure 9: Consensus backbone: radial broadcast structure.

Radial layout of the consensus backbone. Countries are arranged by distance from the central spillover sources, revealing the concentric layers of spillover propagation. Inner-ring countries receive spillovers directly; outer-ring countries are reached through intermediaries.



Source: Author's calculations.

Table 2: Node-level backbone statistics.

Country	In-degree	Out-degree	In-strength	Out-strength	Backbone centrality
ZAF	1	4	0.0759	0.2187	2
COL	3	4	0.1505	0.2027	2
PER	3	3	0.1586	0.1559	2
TKY	1	2	0.0740	0.1223	1
CHL	3	2	0.1524	0.1054	1
PAN	3	2	0.1482	0.1005	1
BRA	3	2	0.1450	0.1002	2
MEX	3	2	0.1540	0.0976	2
PHP	1	2	0.0452	0.0965	1
IDN	1	1	0.0446	0.0452	1
MYS	1	1	0.0470	0.0449	1
CHN	1	1	0.0495	0.0443	1
THA	1	1	0.0449	0.0425	1
KOR	1	0	0.0443	0.0000	0
VNM	1	0	0.0425	0.0000	0
AUS	0	0	0.0000	0.0000	0
AUT	0	0	0.0000	0.0000	0
BEL	0	0	0.0000	0.0000	0
DEN	0	0	0.0000	0.0000	0
DEU	0	0	0.0000	0.0000	0
FRA	0	0	0.0000	0.0000	0
GRB	0	0	0.0000	0.0000	0
HKG	0	0	0.0000	0.0000	0
HUN	0	0	0.0000	0.0000	0
IRL	0	0	0.0000	0.0000	0
ISR	0	0	0.0000	0.0000	0
ITA	0	0	0.0000	0.0000	0
JAP	0	0	0.0000	0.0000	0
NDL	0	0	0.0000	0.0000	0
NOR	0	0	0.0000	0.0000	0
NZD	0	0	0.0000	0.0000	0
POL	0	0	0.0000	0.0000	0
POR	0	0	0.0000	0.0000	0
SPA	0	0	0.0000	0.0000	0
SWE	0	0	0.0000	0.0000	0
USA	0	0	0.0000	0.0000	0

Notes: Statistics computed from the OLS VAR(8) Sinkhorn-normalized consensus backbone at $k = k^* = 3$. “Backbone centrality” is the number of arborescences in which the node appears as a non-leaf (i.e., has at least one outgoing edge in the arborescence). Countries sorted by out-strength in descending order. The theoretical maximum of backbone centrality is $k_r \times k^* = 3 \times 3 = 9$ (three roots, three arborescences each); the empirical maximum of 2 indicates that no single node consistently serves as a non-leaf across all arborescences from all roots, reflecting the diversification goal of the branching-pack construction. Countries with zero in-degree and zero out-degree are isolated from all arborescences; these are predominantly advanced economies whose GFEVD rows are too diffuse to dominate any arborescence path.

Source: Author’s calculations.

In the present sample with the OLS VAR(8) baseline, 13 emerging market economies appear as active transmitters in the backbone (positive out-degree: ZAF, COL, PER, TKY, CHL, PAN, BRA, MEX, PHP, IDN, MYS, CHN, THA). Two further countries appear only as receivers (positive in-degree, zero out-degree: KOR and VNM). The remaining 21 countries are complete isolates not reached by any arborescence, comprising the major advanced economies (AUS, AUT, BEL, DEN, DEU, FRA, GRB, HKG, HUN, IRL, ISR, ITA, JAP, NDL, NOR, NZD, POL, POR, SPA, SWE, USA). This pattern reflects the structure of the VAR(8) GFEVD at the 52-week horizon: advanced-economy sovereign CDS spreads exhibit diffuse, nearly symmetric variance contributions that do not dominate any single arborescence path under Sinkhorn normalization, while the emerging market economies exhibit more concentrated directional spillover channels that the algorithm isolates as the dominant backbone structure.

Finally, as an illustrative summary of the node-level statistics, Table 3 applies a simple quartile-based taxonomy to assign countries to one of three roles: transmitters (top quartile out-strength and bottom quartile in-degree), receivers (top quartile in-degree and bottom quartile out-degree), and bridges (structural articulation points whose removal disconnects the backbone graph). These labels are descriptive summaries of backbone topology, not causal assertions; they should be read as exploratory indicators of relative prominence in the estimated predictive dependence network.

Table 3: Country classification: transmitters, receivers, and bridges.

Role	Countries
Transmitters	None
Receivers	KOR, VNM
Bridges	CHN, IDN, MYS, PHP, THA, ZAF

Notes: Classification based on the OLS VAR(8) Sinkhorn-normalized consensus backbone at $k^* = 3$. Transmitters are defined as countries in the top quartile of out-strength *and* bottom quartile of in-degree; receivers are in the top quartile of in-degree *and* bottom quartile of out-degree; bridges are articulation points of the undirected backbone graph. No country simultaneously satisfies both the top-quartile out-strength and bottom-quartile in-degree conditions, reflecting the concentrated but mutually interconnected nature of the backbone: the top transmitters (ZAF, COL, PER) also receive edges from each other. Bridges (CHN, IDN, MYS, PHP, THA, ZAF) are structural intermediaries whose removal would disconnect the backbone; their presence reflects the relay role of East and Southeast Asian sovereigns in channeling spillovers from Latin American root nodes to broader emerging market receivers.

Source: Author's calculations.

Taken together, the node statistics and role classification show that the baseline backbone is not organized around a single dominant sovereign. Instead, the hierarchy is a compact emerging-market structure in which South Africa, Colombia, and Peru have the largest backbone out-strengths, while several Asian sovereigns serve as intermediaries. The absence of pure transmitters under the quartile rule is informative: the countries with the highest out-strength also receive edges from other active nodes, so their role is closer to that of mutually connected sources than exogenous originators. Korea and Vietnam, by contrast, appear as terminal receivers in the backbone, indicating that their aggregate Diebold–Yilmaz net transmission does not translate into dominant outgoing arborescence paths under Sinkhorn normalization. This distinction is important because the backbone ranks countries by their position in a directed dependence hierarchy, not by their total connectedness in the dense GFEVD matrix.

4.6 Benchmarking against conventional approaches

We now compare the consensus backbone with three conventional approaches: the undirected minimum spanning tree, naive edge-weight thresholding, and the standard Diebold–Yilmaz net directional ranking. Each comparison isolates a different dimension along which the backbone adds informational value.

4.6.1 Comparison with undirected MST

To highlight the value of the directed consensus backbone relative to existing undirected methods, we compute the minimum spanning tree (MST) of the symmetrized variance decomposition matrix, where the weight of undirected edge $\{i, j\}$ is set to $w_{ij} + w_{ji}$, and compare it to the consensus backbone.


Figure 10 makes the effect of symmetrization visible. The consensus backbone in the left panel is concentrated in a small set of directed entries among the active emerging-market sovereigns. Some economically important pairs appear in both directions, such as Turkey–South Africa and Chile–Peru, but they enter as two directed spillover paths rather than as a single symmetric tie. Other entries, such as South Africa to Indonesia or China to Korea, are one-way links that help connect the Latin American to Asian receivers. The sparsity the left panel shows does not reflect loss of information; it is actually the relevant information the backbone extracts: a directed hierarchy over the countries whose normalized spillovers are sufficiently concentrated to dominate arborescence paths.

The MST shown in the right panel answers a different question. Because it is computed on the symmetrized matrix and is forced to span all 36 countries, it necessarily assigns at least one undirected connection to every country, including advanced economies that are isolated in the directed backbone. This produces greater visual coverage and higher retained weight, but it also combines w_{ij} and w_{ji} into a single bilateral relation. Strong MST links such as Hungary–Poland, Hong Kong–New Zealand, Portugal–Spain, Italy–Spain, Israel–United States, and Australia–New Zealand are informative about bilateral closeness in the symmetrized network, but they no longer identify the direction of risk transmission. The comparison should therefore not be read as the MST outperforming the backbone on coverage; the two panels answer distinct questions. The

MST asks which undirected ties connect all countries with high bilateral weight, whereas the consensus backbone asks which directed spillover paths survive normalization and arborescence constraints.

Figure 10: Consensus backbone vs. minimum spanning tree.

Left: directed consensus backbone (OLS VAR(8), Sinkhorn normalization, 27 directed edges). Right: undirected MST of the symmetrized network (35 undirected edges). The MST loses directionality information and represents only the single strongest bilateral path between each country pair, whereas the backbone preserves spillover direction and multiple alternative routes.



Source: Author's calculations.

Table 4 formalizes the comparison. The MST, by construction, discards directionality and retains only $N - 1 = 35$ undirected edges spanning all 36 nodes, yielding a single tree structure with trivially perfect reachability once each edge is treated as bidirectional. The consensus backbone retains 27 directed edges concentrated among 15 active nodes; the MST, drawing on the symmetrized matrix which aggregates both directions of the GFEVD, achieves substantially higher weight coverage (23.8%) because it can freely draw on bidirectional flows. This comparison underscores the different objectives of the two methods: the MST maximizes bilateral connection weight across all node pairs, while the backbone extracts dominant *directional* paths among the countries that exhibit concentrated spillover hierarchies. The backbone reveals not only *which* countries are linked but *who transmits to whom* and through which alternative directed channels.

Table 4: Consensus backbone vs. undirected MST.

Method	Weight coverage (%)	Reachability	Connected components
Consensus backbone	3.8	0.058	1 [‡]
Undirected MST	23.8	1.000 [†]	1

Notes: The consensus backbone retains 27 directed edges among 15 active nodes; the MST retains $N - 1 = 35$ undirected edges covering all 36 nodes. Weight coverage is the ratio of retained edge weight to total directed network weight; MST weight coverage uses the symmetrized (bidirectional) edge weights. Reachability is the fraction of ordered pairs (i, j) across all 36 nodes for which a directed path exists; the backbone’s low reachability (5.8%) reflects the 21 isolated advanced-economy nodes and the directed (non-spanning) nature of the arborescences. [†] Reachability equals 1 by construction: the MST is a spanning tree covering all nodes bidirectionally. [‡] The backbone graph contains one weakly connected component among the 15 active nodes; the 21 isolated countries are not part of the backbone graph.

Source: Author’s calculations.

Undirected sparsification benchmarks such as the disparity filter and salience skeleton are not directly comparable here because they are designed for symmetric graphs and would require a preliminary symmetrization step that discards spillover direction. For that reason, the more meaningful directed benchmarks in this paper are the standard Diebold–Yilmaz rankings and the directed VAR robustness checks reported below.

4.6.2 Comparison with naive thresholding

An alternative sparsification approach is to retain the top- N edges by weight, where N matches the number of edges in the consensus backbone. Table 5 compares the consensus backbone with this naive thresholding approach at the same edge density.

Table 5: Consensus backbone vs. naive edge-weight thresholding.

Method	Weight coverage (%)	Reachability	Connected components
Consensus backbone	3.8	0.058	1 [‡]
Top-27 threshold	6.3	0.037	23

Notes: Both methods retain 27 edges. Reachability is the fraction of ordered node pairs (i, j) across all 36 nodes for which a directed path exists from i to j . [‡] The backbone graph contains one weakly connected component among the 15 active nodes; the 21 isolated countries form singleton components if counted separately, giving 22 total components under that convention. The top-27 threshold graph fragments into 23 components, confirming that weight-maximizing selection does not preserve connectivity. Both reachability figures are low because the 21 isolated advanced-economy nodes contribute unreachable pairs regardless of the sparsification method; the key structural comparison is that the backbone achieves higher reachability than threshold among the active nodes it covers.

Source: Author’s calculations.

The comparison illustrates the trade-off between local weight maximization and path preservation. The top-27 threshold graph retains more aggregate weight than the consensus backbone (6.3 percent versus 3.8 percent), but it fragments into 23 connected components and achieves lower

reachability. This occurs because thresholding evaluates each edge independently: it can select many high-weight links concentrated around the same countries without ensuring that those links compose into directed spillover paths. The consensus backbone imposes arborescence structure, so it retains a lower-weight edge set that nevertheless forms one weakly connected component among the active countries and preserves more directed paths. The gain from the backbone is therefore not higher retained weight, but a more interpretable sparse topology.

4.6.3 Comparison with standard Diebold–Yilmaz rankings

To assess whether the consensus backbone uncovers structure beyond what standard connectedness measures already reveal, Table 6 compares backbone-based country importance (out-strength rank) with the standard Diebold–Yilmaz net directional spillover ranking.

Table 6: Diebold–Yilmaz net spillover vs. backbone out-strength (top 15).

Country	DY net spillover	DY rank	Backbone out-strength	BB rank
IDN	0.4620	1	0.0452	10
PHP	0.4580	2	0.0965	9
MYS	0.4085	3	0.0449	11
COL	0.3147	4	0.2027	2
MEX	0.3115	5	0.0976	8
PAN	0.3034	6	0.1005	6
CHN	0.2817	7	0.0443	12
THA	0.2610	8	0.0425	13
CHL	0.2449	9	0.1054	5
KOR	0.2442	10	0.0000	14
PER	0.2336	11	0.1559	3
VNM	0.2322	12	0.0000	14
BRA	0.1810	13	0.1002	7
BEL	0.1532	14	0.0000	14
IRL	0.0686	15	0.0000	14

Notes: DY net spillover = spillover transmitted to others – spillover received from others, computed from the OLS VAR(8) GFEVD. Backbone out-strength from the Sinkhorn-normalized consensus backbone at $k^* = 3$. Countries with zero backbone out-strength share the same rank (tied at rank 14). Spearman rank correlation between the two rankings across all 36 countries: $\rho = 0.64$, $p < 0.001$.

Source: Author’s calculations.

The two rankings broadly agree on the East Asian and Latin American countries as dominant transmitters under the VAR(8) baseline, but diverge for others. Indonesia, Philippines, and Malaysia top the DY net rankings but have low backbone out-strength, reflecting spillovers that are diffused broadly rather than concentrated along dominant directed paths. Conversely, Colombia, Peru, and South Africa achieve the highest backbone out-strength while ranking moderately in DY net terms, because the backbone isolates the directional concentration of their spillovers rather than their aggregate net balance. Korea and Vietnam rank in the top-12

of DY net transmitters but have zero backbone out-strength since their influence is measured by the DY framework as high aggregate net transmission but is too diffuse to dominate any arborescence path. The Spearman rank correlation across all 36 countries is $\rho = 0.64$ ($p < 0.001$), a moderate positive association that confirms the backbone captures related but structurally distinct information from the standard DY rankings.

4.7 Robustness and sensitivity analysis

This section examines the stability of the consensus backbone to changes in estimation method, normalization, forecast horizon, and root selection.

4.7.1 Sensitivity to VAR lag order

The baseline uses OLS VAR(8). To assess sensitivity to this choice, Table 7 reports backbone statistics for lag orders of 1, 4, 8, and 12—spanning the feasible range from the IC-minimum (lag 1) to near the feasibility ceiling (lag 13). For each specification we compute the OLS GFEVD at $H = 52$ weeks, apply the weighted composite criterion to select k^* , extract the Sinkhorn backbone, and compare the resulting edge set and node rankings against the VAR(8) baseline using the Jaccard similarity of edge sets and the Spearman rank correlation of out-strength.

Table 7: Sensitivity of the consensus backbone to VAR lag order.

Specification	k^*	Edges	Active nodes	Weight cov. (%)	Jaccard	ρ_S
VAR(1)	5	21	8	3.3	0.455	0.821
VAR(4)	5	40	15	5.6	0.396	0.910
VAR(8) [baseline]	3	27	15	3.8	1.000	1.000
VAR(12)	4	28	13	3.3	0.222	0.611

Notes: All specifications use the same 36-country, 543-observation weekly sample (2015–2025), Sinkhorn normalization, and the same backbone hyperparameters ($k_r = 3$ roots, $k_p = 5$ arborescences per root, penalty $\gamma = 0.6$). The GFEVD horizon is 52 weeks in all cases. “Jaccard” is the Jaccard similarity of the edge set relative to the VAR(8) baseline; “ ρ_S ” is the Spearman rank correlation of backbone out-strength relative to VAR(8). VAR(12) is included for comparison but has spectral radius 0.984 (260 companion eigenvalues in the range $[0.95, 1.0)$), suggesting near-non-stationarity that may distort long-horizon GFEVD.

Source: Author’s calculations.

The backbone is moderately stable across lag orders of 4–8, which span the standard range used in the weekly DY literature. The VAR(4) and VAR(8) backbones share the same 15 active nodes and a Spearman rank correlation of 0.91, indicating that the dominant spillover hierarchy is consistent across these specifications. The VAR(1) backbone is sparser (21 edges, 8 active nodes) because a single-lag GFEVD at the 52-week horizon is dominated by the instantaneous covariance structure and misses medium-run dynamics. The VAR(12) backbone diverges more substantially (Jaccard = 0.22, $\rho_S = 0.61$) and rests on a near-nonstationary system (spectral radius 0.984), making it the least reliable specification. These results support the VAR(8) choice as a robust intermediate that captures meaningful spillover dynamics without the stability concerns of longer lags.

4.7.2 Bootstrap edge-stability analysis


The backbone is a nonlinear function of the estimated GFEVD matrix, which is itself subject to sampling uncertainty. To assess which backbone edges are most stable, we conduct a block bootstrap over the full sample: the VAR(8) residuals are resampled in non-overlapping blocks of 13 weeks (one quarter), the VAR is re-estimated on each pseudo-sample, a new GFEVD is computed, and the Sinkhorn backbone is re-extracted. This is repeated for 200 replicates.

Figure 11 reports the selection frequency of every directed edge across replicates; blue outlines mark the edges present in the baseline backbone.

The results indicate substantial sampling uncertainty in individual edge identities, which is consistent with the well-known sensitivity of large-dimensional VARs to sampling variation. Blue outlines in Figure 11 mark the 27 edges of the OLS VAR(8) baseline backbone; the selection frequencies for these edges vary considerably across replicates. Rather than undermining the method, this finding underscores that the backbone should be interpreted as a conditional, exploratory summary of the estimated GFEVD—a tool for prioritizing structure in a dense network—rather than as a set of precisely identified bilateral relationships. What appears to be more stable is the coarse hierarchy of important nodes and dominant channels, not the literal identity of each bilateral edge. A full rolling-window bootstrap that also tracks stability of the selected root set and of k^* across windows is left for future work.

Figure 11: Bootstrap edge-selection frequencies.

Each cell shows the fraction of 200 block-bootstrap replicates (block size = 13 weeks, $\text{VAR}(8)$) in which the corresponding directed edge appears in the Sinkhorn-normalized backbone. Blue outlines indicate edges present in the baseline backbone. Warmer colours indicate higher selection frequency.



Source: Author's calculations.

4.7.3 Sensitivity to the GFEVD horizon

To check whether the forecast horizon drives the backbone mechanically, we re-estimate a parsimonious OLS $\text{VAR}(8)$ on the full weekly sample and recompute the Sinkhorn backbone at horizons of 4, 12, 26, and 52 weeks. Figure 12 summarizes the main patterns, and Table 8 reports the exact values.

Table 8: Horizon sensitivity of the OLS VAR(8) backbone.


H (weeks)	k^*	Edges	Active nodes	Weight coverage (%)	Reachability	Jaccard vs. $H = 52$	Spearman ρ vs. $H = 52$
4	3	33	15	4.7	0.062	0.765	0.979
12	3	20	8	3.0	0.035	0.741	0.846
26	3	27	15	3.8	0.058	1.000	1.000
52	3	27	15	3.8	0.058	1.000	1.000

Notes: The OLS VAR(8) is estimated on the full weekly sample, and the same Sinkhorn normalization and backbone hyperparameters are used in every horizon. The selected roots are identical across horizons: ZAF, TKY, and COL.

Source: Author's calculations.

Figure 12: Sensitivity of the OLS VAR(8) backbone to the GFEVD horizon.

The top panel shows how the backbone size changes with the GFEVD horizon. The bottom panel compares each horizon-specific backbone with the 52-week benchmark using edge-set Jaccard overlap and Spearman rank correlation of out-strength. The selected roots are unchanged across horizons (ZAF, TKY, and COL), but the edge set becomes sparser at shorter horizons.



Source: Author's calculations.

The horizon sensitivity results suggest that the main backbone hierarchy is not mechanically driven by the choice of a 52-week forecast horizon. The selected roots are identical at all horizons, and the optimal number of arborescences remains $k^* = 3$ throughout. The 26-week and 52-week backbones are identical under these diagnostics, indicating that the dominant long-run spillover structure has already stabilized by roughly six months. At the 4-week horizon, the edge set is somewhat larger, but the active-node count remains 15 and the out-strength ranking is almost unchanged relative to the baseline ($\rho_S = 0.979$). The 12-week horizon is the main exception, producing a smaller active subgraph, although its rank correlation with the 52-week benchmark remains high. Overall, the evidence supports interpreting the baseline hierarchy as a persistent medium- to long-horizon structure rather than an artifact of one arbitrary forecast horizon.

4.7.4 Sensitivity to root selection

Using the same OLS VAR(8) and 52-week horizon, we vary the number of out-strength roots k_r from 1 to 6. The goal is to assess whether the backbone is driven by a narrow seed choice or whether it is stable once the main spillover sources are included. Table 9 summarizes the results, and Figure 13 visualizes the out-strength-root cases.

The root-selection exercise indicates that the baseline backbone is not sensitive to the exact number of out-strength roots once the principal sources are included. With only South Africa as a root, the backbone already recovers much of the baseline structure, but its lower Jaccard overlap and rank correlation show that a single source is too restrictive. Adding Turkey as a second root is sufficient to reproduce the full baseline edge set, and adding further out-strength roots leaves the selected backbone unchanged. This saturation supports the use of $k_r = 3$ as a conservative baseline rather than a finely tuned choice. The in-strength counterfactual, by contrast, produces a much denser backbone over all 36 countries with very low overlap relative to the baseline, confirming that the economic interpretation depends on rooting the arborescences at transmitters rather than receivers.

Table 9: Root-selection sensitivity of the OLS VAR(8) backbone.


Root rule	Roots	k^*	Edges	Active nodes	Weight coverage (%)	Jaccard vs. baseline	Spearman ρ vs. baseline
Out-strength $k_r = 1$	ZAF	2	20	15	2.8	0.741	0.858
Out-strength $k_r = 2$	ZAF, TKY	3	27	15	3.8	1.000	1.000
Out-strength $k_r = 3$	ZAF, TKY, COL	3	27	15	3.8	1.000	1.000
Out-strength $k_r = 4$	ZAF, TKY, COL, CHL	3	27	15	3.8	1.000	1.000
Out-strength $k_r = 5$	ZAF, TKY, COL, CHL, PER	3	27	15	3.8	1.000	1.000
Out-strength $k_r = 6$	ZAF, TKY, COL, CHL, PER, BRA	3	27	15	3.8	1.000	1.000
In-strength $k_r = 3$	HUN, IRL, BRA	2	42	36	7.2	0.190	–

Notes: Jaccard overlap and Spearman rank correlations are computed relative to the baseline out-strength specification with $k_r = 3$. The in-strength counterfactual yields a backbone with very low overlap (Jaccard = 0.19) and a rank correlation that is not computable due to the different active-node sets, reinforcing the interpretation of out-strength as the economically relevant root criterion.

Source: Author's calculations.

Figure 13: Sensitivity of the OLS VAR(8) backbone to root selection.

The top panel shows backbone size as the number of out-strength roots increases. The bottom panel compares each root set with the baseline $k_r = 3$ specification using edge-set Jaccard overlap and Spearman rank correlation of out-strength. Once $k_r \geq 2$, the backbone is unchanged relative to the baseline.



Source: Author's calculations.

5 Discussion: advanced economies and the backbone

Perhaps the most striking empirical finding of the illustrative application is the complete absence of advanced economies from the consensus backbone. Under Sinkhorn normalization, all 21 advanced economies in the sample—including the United States, Germany, Japan, France, and the United Kingdom—appear as complete isolates with zero in-degree, zero out-degree, and zero backbone centrality (Table 2). The entire active backbone consists of 15 emerging market and commodity-linked sovereigns, organized into a hierarchical structure rooted in Latin American and Asian economies. This section explores why this pattern arises and what it reveals about the structure of sovereign risk transmission.

5.1 Why advanced economies drop out

The dropout of advanced economies is not an artifact of the backbone method but a direct consequence of the topology of the Sinkhorn-normalized variance decomposition network. In the raw GFEVD matrix, advanced economies contribute substantial variance shares to one

another, likely reflecting deep financial integration, common factor exposures, and correlated policy cycles. These contributions, however, are spread broadly across bilateral pairs: no single advanced economy dominates another’s forecast error variance by a wide margin. By equalizing row and column sums, Sinkhorn normalization removes these scale effects and leaves the *relative concentration* of spillover channels—edges that are disproportionately strong given the overall emission and absorption profiles of their endpoints.

Emerging market sovereigns, by contrast, exhibit sharper bilateral asymmetries. A country such as South Africa or Colombia may not contribute the largest aggregate variance share to the system, but a few of its outgoing links carry disproportionate influence. After Sinkhorn normalization, these channels become high-weight edges, and the Chu–Liu/Edmonds algorithm selects them into the arborescences. Advanced economies, whose normalized profiles are flatter, fail to dominate any arborescence path and therefore do not appear in the backbone.

The comparison with the standard Diebold–Yilmaz rankings in Table 6 supports this interpretation. Several advanced economies rank among the top net spillover transmitters in the raw DY framework. Belgium and Ireland, for instance, appear in the top 15 by net directional spillover yet have zero backbone presence. The DY net index aggregates all bilateral channels into a single scalar, rewarding countries with high total volume; the backbone instead asks whether a country dominates specific bilateral transmission channels after controlling for scale.

5.2 Implications for sovereign risk transmission

The advanced-economy dropout carries substantive implications for how sovereign risk spillovers are understood and monitored. First, it separates broad connectedness from backbone centrality. Advanced economies can be systemically important in aggregate while contributing little to the sparse transmission skeleton because their influence is diffuse rather than organized through dominant bilateral channels. The centrality of emerging market sovereigns points to a different structure: spillovers concentrated along specific routes, plausibly reflecting regional trade linkages, commodity-price co-exposures, and investor portfolio segmentation.

Second, the result matters for macro-prudential surveillance. Standard aggregate connectedness indices, which often place advanced economies near the top of systemic importance rankings by virtue of their size and integration, identify countries with large total spillover balances. The backbone instead highlights the clearest directed routes through which sovereign stress can propagate. On this criterion, attention shifts toward the emerging market core of the backbone, where transmitter roles are more sharply defined.

Third, the bridge role assigned to several East and Southeast Asian sovereigns acquires additional significance in light of the advanced-economy dropout (Table 3). These countries serve as structural intermediaries connecting Latin American root nodes to broader emerging market receivers. Their bridge status implies that shocks originating in Latin America reach the Asian periphery through specific relay nodes rather than through advanced-economy hubs. This pattern is consistent with a portfolio-rebalancing channel, in which emerging-market dedicated investors reallocate across regions in response to sovereign credit events.

5.3 Robustness of the dropout pattern

The dropout of advanced economies is not confined to the baseline specification. Under the VAR(4) specification (Table 7), the backbone retains the same 15 active emerging market nodes. The pattern is also stable across GFEVD horizons of 4 to 52 weeks (Table 8), with the set of active nodes remaining exclusively emerging market at all horizons except the 12-week specification, where the backbone contracts to 8 active nodes—still entirely emerging market economies. The root-selection sensitivity analysis (Table 9) further confirms that the backbone structure is unchanged once at least two out-strength roots are included, and that all selected roots are emerging market sovereigns regardless of how many roots are specified.

It is important to note that this finding is conditional on the Sinkhorn normalization, which is by design the most aggressive scale equalizer. Under raw weights, the backbone is substantially denser (177 edges) and includes advanced-economy nodes that dominate through sheer magnitude. Under row- or column-stochastic normalization, intermediate patterns emerge. The choice of Sinkhorn as the primary specification is deliberate: it isolates the structural skeleton of spillover concentration, stripped of the scale advantages that large, highly integrated economies enjoy by default. The finding that advanced economies vanish under this normalization is thus not a limitation but a feature—it reveals that their prominence in the raw network is driven by scale rather than by the existence of dominant bilateral transmission channels.

5.4 Caveats and directions for further analysis

Several caveats apply. The backbone is extracted from a single GFEVD estimated over a fixed sample period; advanced economies might re-enter the backbone during specific episodes of concentrated bilateral stress, such as the European sovereign debt crisis, when peripheral-to-core spillover channels became highly asymmetric. A rolling-window backbone analysis would clarify whether the dropout is a permanent structural feature or an artifact of the full-sample averaging that smooths out episodic concentration.

Additionally, the GFEVD captures predictive variance shares, not structural causation. The absence of advanced economies from the backbone does not imply that they are unimportant for global sovereign risk: it implies that their importance operates through diffuse, multi-channel mechanisms that the backbone, by construction, does not retain. Complementary methods, such as factor-augmented VAR models that explicitly account for common factors, could be used alongside the backbone to provide a more complete picture of sovereign risk interdependence.

6 Conclusion

This paper has proposed the consensus backbone, constructed as the union of branching packs of maximum arborescences, as a principled method for extracting sparse, directed, and interpretable structures from dense economic and financial networks. The method contributes a general framework for summarizing dense directed dependence with a sparse backbone that preserves the dominant directional structure of the network. The Diebold–Yilmaz spillover application shows how the method works in practice on a 36-country example, where the backbone retains a

non-trivial share of total network weight (3.8%) using only 2.1% of edges under the OLS VAR(8) specification.

Beyond the methodological contribution, the application yields a substantive empirical finding: under Sinkhorn normalization, all 21 advanced economies in the sample are absent from the consensus backbone, while the entire active spillover hierarchy consists of 15 emerging market and commodity-linked sovereigns. As discussed in Section 5, this pattern reveals a fundamental asymmetry in sovereign risk transmission. Advanced-economy spillovers propagate through diffuse, nearly symmetric channels that do not lend themselves to hierarchical representation, whereas emerging market spillovers are concentrated along specific bilateral pathways that the backbone isolates as the dominant directional structure. This finding suggests that standard aggregate connectedness indices, which place advanced economies at the top of systemic importance rankings by virtue of their scale, may obscure the distinct topology of sovereign risk transmission across the two groups.

The method identifies three illustrative structural roles—transmitters, receivers, and bridges—in the dominant predictive dependence skeleton of the example network. Transmitter countries occupy the roots of the backbone’s arborescences, reflecting high relative out-strength in the Sinkhorn-normalized GFEVD conditional on the estimated VAR. Bridge countries are structural intermediaries whose removal would disconnect the backbone graph; as exploratory indicators, they may point to natural targets for cross-border risk monitoring, pending validation against observable bilateral exposure data. The diagnostic framework for selecting k^* ensures that the backbone is both parsimonious and representative, with diminishing returns providing a natural stopping rule.

More broadly, the consensus backbone extends tree- and hierarchy-based approaches to directed economic and financial networks by preserving multiple dominant arborescent structures rather than forcing the network into a single tree. This makes the method useful beyond the sovereign-spillover application in settings where policy analysis requires sparse, interpretable summaries of dense exposure or dependence networks.

An important avenue for applied work using the consensus backbone is external economic validation: linking backbone edges to observable bilateral channels such as BIS cross-border bank claims, trade exposures, portfolio investment linkages, or common dependence on US dollar funding conditions would help establish whether the dominant spillover paths identified by the method correspond to economically meaningful spillover channels.

Several extensions merit future investigation. First, *time-varying consensus backbones* could be constructed by rolling the VAR estimation window and tracking how the backbone evolves across regimes. A rolling-window analysis, combining bootstrap-based edge-stability tests (Section 4.7.2) with structural break detection, would clarify whether the advanced-economy dropout is a permanent structural feature or whether concentrated bilateral channels among advanced economies emerge during specific stress episodes, such as the European sovereign debt crisis. Second, *community detection on the backbone*—for instance, applying the Louvain algorithm to the backbone graph—could reveal whether sovereign risk spillover pathways cluster

by region, funding currency, or other shared economic affinities. Such a decomposition would complement the transmitter, receiver, and bridge roles highlighted above, and would further clarify how the consensus backbone organizes dense directional dependence into a sparse and interpretable structure.

References

- Anand, Kartik, Iman van Lelyveld, Ádám Banai, Soeren Friedrich, Rodney Garratt, Grzegorz Halaj, Jose Figue, et al. 2018. “The missing links: A global study on uncovering financial network structures from partial data.” *Journal of Financial Stability* 35:107–119. ISSN: 1572-3089. <https://doi.org/https://doi.org/10.1016/j.jfs.2017.05.012>.
- Battiston, Stefano, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli. 2012. “DebtRank: Too central to fail? Financial networks, the FED and systemic risk.” *Sci. Rep.* 2:541. <https://doi.org/10.1038/srep00541>.
- Bech, Morten L., and Kim Sørensen. 2018. “The topology of the Danish interbank market.” *Danmarks Nationalbank Working Paper*.
- Bérczi, Kristóf, and András Frank. 2016. “Supermodularity in Unweighted Graph Optimization I: Branchings and Matchings.” *arXiv preprint*, arXiv: 1608.05722 [math.CO]. <https://arxiv.org/abs/1608.05722>.
- Bougheas, Spiros, and Alan Kirman. 2014. Complex Financial Networks and Systemic Risk: A Review. Technical report. CESifo Working Paper No. 4756. https://www.econstor.eu/bitstream/10419/96832/1/cesifo_wp4756.pdf.
- Camerini, Paolo M., Luigi Fratta, and Francesco Maffioli. 1979. “A Note on Finding Optimum Branchings.” *Networks* 9 (4): 309–312. <https://doi.org/10.1002/net.3230090403>.
- Chakrabarty, Deeparnab, and Sanjeev Khanna. 2018. “Better and Simpler Error Analysis of the Sinkhorn–Knopp Algorithm for Matrix Scaling.” In *Proceedings of the Symposium on Simplicity in Algorithms (SOSA)*, 4:1–4:11. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. <https://doi.org/10.4230/OASIcs.SOSA.2018.4>.
- Chu, Y. J., and T. H. Liu. 1965. “On the Shortest Arborescence of a Directed Graph.” *Scientia Sinica* 14:1396–1400.
- Coscia, Michele. 2018. “Using arborescences to estimate hierarchicalness in directed complex networks.” *PLoS ONE* 13 (1): e0190825. <https://doi.org/10.1371/journal.pone.0190825>.
- Diebold, Francis X., and Kamil Yilmaz. 2012. “Better to Give than to Receive: Predictive Directional Measurements of Volatility Spillovers.” *International Journal of Forecasting* 28 (1): 57–66. <https://doi.org/10.1016/j.ijforecast.2011.02.006>.
- Edmonds, Jack. 1967. “Optimum Branchings.” *Journal of Research of the National Bureau of Standards. Section B, Mathematics and Mathematical Physics* 71B (4): 233–240. https://nvlpubs.nist.gov/nistpubs/jres/71b/jresv71bn4p233_a1b.pdf.
- Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson. 2019. “Networks and Economic Policy.” *Oxford Review of Economic Policy* 35 (4): 565–614. <https://doi.org/10.1093/oxrep/grz024>.

- Emmert-Streib, Frank, Ahmed Musa, Karolis Baltakys, Juho Kanninen, Shailesh Tripathi, Olli Yli-Harja, Herbert Jodlbauer, and Matthias Dehmer. 2018. “Computational Analysis of the Structural Properties of Economic and Financial Networks.” *Journal of Network Theory in Finance* 4:1–32. <https://doi.org/10.21314/JNTF.2018.043>.
- Esmalifalak, Hamid. 2021. “Correlation networks in economics and finance: A review of methodologies and innovations.” *Journal of Economic Surveys* 35:1200–1232. <https://doi.org/10.1111/joes.12655>.
- Gabauer, David, Ioannis Chatziantoniou, and Alexis Stenfors. 2023. “Model-free connectedness measures.” *Finance Research Letters* 54 (103804). <https://doi.org/10.1016/j.frl.2023.103804>.
- Gaffeo, Edoardo, and Antonio Molinari. 2016. A functional perspective on financial networks. Technical report. SSRN Working Paper.
- Georgiadis, Leonidas. 2003. “Arborescence Optimization Problems Solvable by Edmonds’ Algorithm.” *Theoretical Computer Science* 301 (1-3): 427–437. [https://doi.org/10.1016/S0304-3975\(02\)00888-5](https://doi.org/10.1016/S0304-3975(02)00888-5).
- Grady, Daniel, Christian Thiemann, and Dirk Brockmann. 2012. “Robust classification of salient links in complex networks.” *Nature Communications* 3 (864). <https://doi.org/10.1038/ncomms1847>.
- Huremovic, Kenan, Gabriel Jiménez, Enrique Moral-Benito, José-Luis Peydró, and Fernando Vega-Redondo. 2020. Production and Financial Networks in Interplay: Crisis Evidence from Supplier–Customer and Credit Registers. Technical report. Barcelona School of Economics Working Paper No. 1191 / CEPR Discussion Paper No. DP15277. https://www.econstor.eu/bitstream/10419/222281/1/FinShocksPN_Final_July9.pdf.
- Idel, Martin. 2016. “A Review of Matrix Scaling and Sinkhorn’s Normal Form for Matrices and Positive Maps.” *arXiv preprint*, arXiv: 1609.06349 [math.FA]. <https://arxiv.org/abs/1609.06349>.
- Joseph, Andreas, Irena Vodenska, Eugene Stanley, and Guanrong Chen. 2014. Netconomics: Novel Forecasting Techniques from the Combination of Big Data, Network Science and Economics. Technical report. arXiv preprint arXiv:1403.0848. <https://arxiv.org/abs/1403.0848>.
- Knight, Philip A. 2008. “The Sinkhorn–Knopp Algorithm: Convergence and Applications.” *SIAM Review* 50 (1): 57–76. <https://doi.org/10.1137/060659624>.
- Mantegna, Rosario N. 1999. “Hierarchical structure in financial markets.” *The European Physical Journal B - Condensed Matter and Complex Systems* 11 (1): 193–197. <https://doi.org/10.1007/s100510050929>.
- Mastrandrea, Rossana, Tiziano Squartini, Giorgio Fagiolo, and Diego Garlaschelli. 2014. “Reconstructing the world trade multiplex: the role of intensive and extensive biases.” *Phys. Rev. E* 90:062804. <https://doi.org/10.1103/PhysRevE.90.062804>.

- Meier, Julia, Prejaas Tewarie, and Piet Van Mieghem. 2015. “The Union of Shortest Path Trees of Functional Brain Networks.” *Brain Connectivity* 5 (9): 575–581. <https://doi.org/10.1089/brain.2015.0349>.
- Pesaran, H. Hashem, and Yongcheol Shin. 1998. “Generalized impulse response analysis in linear multivariate models.” *Economic Letters* 58 (1): 17–29. [https://doi.org/10.1016/S0165-1765\(97\)00214-0](https://doi.org/10.1016/S0165-1765(97)00214-0).
- Rozum, Jordan C., and Luis Rocha. 2024. “The Ultrametric Backbone is the Union of all Minimum Spanning Forests.” *Journal of Physics: Complexity* 5 (3). <https://doi.org/10.1088/2632-072X/ad679e>.
- Serri, Matteo, Stefano Battiston, Guido Caldarelli, and Diego Garlaschelli. 2016. “Interbank contagion and the probability of cascade failures.” *Phys. Rev. E* 94:062312. <https://doi.org/10.1103/PhysRevE.94.062312>.
- Sharma, Smriti, Anindya Chakrabarti, and Anirban Chakraborti. 2018. Multi-layered Network Structure: Relationship Between Financial and Macroeconomic Dynamics. Technical report. arXiv preprint arXiv:1805.06829. <https://arxiv.org/abs/1805.06829>.
- Sinkhorn, Richard. 1967. “Diagonal Equivalence to Matrices with Prescribed Row and Column Sums.” *The American Mathematical Monthly* 74 (4): 402–405. <https://doi.org/10.2307/2314570>.
- Spelta, Alessandro, and Tanya Araújo. 2011. “The topology of cross-border exposures: Beyond the minimal spanning tree approach.” *Physica A: Statistical Mechanics and its Applications* 390 (23–24): 4368–4379. <https://doi.org/10.1016/j.physa.2011.06.071>.
- Tarjan, Robert E. 1977. “Finding Optimum Branchings.” *Networks* 7:25–35. <https://doi.org/10.1002/net.3230070103>.
- Valle, Mauricio A., and Felipe Urbina. 2022. “The Backbone of the Financial Interaction Network Using a Maximum Entropy Distribution.” *Advances in Complex Systems* 25 (4): 2250006. <https://doi.org/10.1142/S0219525922500060>.
- Vitali, Stefania, James B. Glattfelder, and Stefano Battiston. 2011. “The network of global corporate control.” *PLoS ONE* 6 (10): e25995. <https://doi.org/10.1371/journal.pone.0025995>.
- Wang, Zhaojun, Jun Zhang, Shidong Zhang, and Zhiyuan Zheng. 2023. “Efficient and Effective Directed Minimum Spanning Tree.” *Mathematics* 11 (9): 2200. <https://doi.org/10.3390/math11092200>.
- Yassin, Abdelkader, Abderrahim Haidar, Hocine Cherifi, Hachemi Seba, and Olivier Togni. 2023. “An evaluation tool for backbone extraction techniques in weighted complex networks.” *Scientific Reports* 13 (17000): 1–13. <https://doi.org/10.1038/s41598-023-42076-3>.